

A Qualitative Investigation of Selected Infrared Flow Visualization Image Processing Techniques

Theodore J. Garbeff II^{*}, Jennifer K. Baerny[†]
NASA Ames Research Center, Moffett Field, Ca, 94035

The deployment and integration of high-sensitivity infrared cameras in a transonic wind tunnel test environment has resulted in a unique capability to image aerodynamic phenomena in real-time. Multi-camera infrared flow visualization data systems are now routinely utilized at the NASA Ames Unitary Plan Wind Tunnel. The small flow-induced temperature gradients on the surface of the wind tunnel test article coupled with the high bit-depth of the infrared camera sensor makes the processing of the image data critically important. An image processing routine must enhance features of interest with minimal artifacts. Additionally, the production wind tunnel test environment demands that these processed images are made available in a real-time, automatic fashion. Therefore, any image processing routine must be computationally economical and enhance the image data with minimal input from a human operator. The following seeks to qualitatively explore selected image processing techniques by assessing their effectiveness to resolve flow features on a wind tunnel test article. A multi-scale contrast enhancement technique is discussed as well as a new implementation of a multi-scale, non-interpolated adaptive histogram equalization. Finally, a novel method is introduced that demonstrates the ability to resolve flow features imaged on bare-steel test articles possessing low emissivity. This method makes use of dynamic mode decomposition and discrete-time filtering to separate the background reflections that dominate low emissivity surfaces from the aerodynamic driven surface temperature gradients. This process will be shown to resolve the onset of boundary layer transition on a bare metal wing as well as identify and resolve hidden features in the image data. While the implementation of this technique is very preliminary it demonstrates the potential to extend the application of infrared flow-visualization within the wind tunnel test environment.

I. Introduction

The deployment and integration of high-sensitivity infrared (IR) cameras in a transonic wind tunnel test environment has resulted in a unique capability to image aerodynamic phenomena in real-time. Multi-camera IR flow visualization data systems are now routinely utilized [1] at the NASA Ames Unitary Plan Wind Tunnel (UPWT). The small flow-induced temperature gradients on the surface of the wind tunnel test article coupled with the high bit-depth of the IR camera sensor makes the processing of the image data critically important. In order to successfully perform IR flow visualization on a wind tunnel test article, it is important that the model surfaces being imaged possess high-emissivity. This is critical as the success of this technique relies on identifying temperature gradients occurring on the surface of the test article. A low emissivity surface results in the area of interest being dominated by reflections of the surrounding test environment. This is compounded by the fact that the longer wavelengths of mid-wave (3-5um) IR are reflected more readily than visible wavelengths [2], resulting in a “shinier” surface than intuition would suggest. Given the high-strength requirements of wind tunnel test articles, models are often constructed of stainless steel or aluminum. An aluminum or stainless-steel model with a surface finish on the order of 25 micro-inches likely has surface emissivities of under 10% in the mid-wave IR band. When painted with a lusterless coating and sanded to an equivalent surface finish emissivities of 85% are attainable. Therefore, when conducting IR flow visualization, the wind tunnel test article is often painted with a high-emissivity coating towards the end of the model preparation process. Care must be taken to ensure the paint is applied evenly with a smooth surface finish. It is possible to achieve a surface finish capable of supporting laminar flow but there are some model geometries that are extremely sensitive to outer mold line and subsequently precludes the article from being painted.

In the case of IR imaging on a wind tunnel model surface of high-emissivity, the image data is generally processed in a manner that seeks to enhance contrast. In section A we discuss methods of localized contrast enhancement that have worked well at the NASA Ames UPWT. The case of IR imaging on low-emissivity or unpainted surfaces is challenging and contrast enhancement alone is insufficient in resolving flow driving features from background reflections. If one assumes that aerodynamic flow features occur at different time scales than the background reflection of the tunnel, it may be possible to distinguish one from the other. In section B we will present a method of IR video processing that seeks to separate spatiotemporal phenomena by utilizing dynamic mode decomposition (DMD) and resolve aerodynamic flow features on bare metal wind tunnel models.

^{*} Aerospace Engineer, Wind Tunnel Systems Branch, MS 227, Member AIAA

[†] Aerospace Engineer, Wind Tunnel Systems Branch, MS 227

II. Methods

A. Contrast Enhancement

The following discussion uses a reference set of images to compare various image processing approaches. This reference image set is of a wind tunnel model in the 11-by-11 foot test section of the NASA Ames UPWT at a Mach number of 0.95. This test article has been painted with a high-emissivity (~ 0.88) coating for the purpose of IR flow visualization. The imagery was acquired by a FLIR SC8203 mid-wave IR camera whose sensor possess a depth of 14-bits per pixel. To view an image through most human machine interfaces, it is necessary to rescale each pixel from 14-bits to 8-bits. Upon examination of the histogram (Figure 1) of a raw image it is apparent that this is not a straight forward task. Typically, the histograms of the unprocessed, 14-bit images are not normally distributed and contain a wide dynamic range. A flow feature driven temperature gradient may only be on the order of tens of counts and consequently will be hidden within the large dynamic range of the image. Classically, the task of enhancing features within an image falls under an image processing subset known as contrast enhancement [3]. Often times this results in an image whose histogram bins have been more evenly distributed across the display range. An image processing routine must enhance features of interest with minimal artifacts. Additionally, the production wind tunnel test environment demands that these processed images are made available in a real-time, automatic fashion [1,4]. Therefore, any image processing routine must be computationally economical and enhance the image data with minimal input from a human operator. The following seeks to qualitatively explore selected image processing techniques by assessing their effectiveness to resolve flow features on a wind tunnel test article. A multi-scale contrast enhancement technique is introduced and compared with a non-interpolated adaptive histogram equalization method.




Figure 1: Raw IR image histogram (14-bit Pixel Depth)

Figure 2 depicts an 8-bit linearly rescaled image (left) whose original histogram is detailed in Figure 1. In the following discussion this image will be used as a reference to compare image processing techniques. The rescaling (left) performed is linear so that the shape of the original histogram is preserved. The outline of high contrast features are visible. These features are mainly related to model geometry (wings and tail), as well as the wind tunnel test section slotted walls. Some hints of flow driven features can be seen near the leading edges on the surface of the wings, however the lack of contrast makes them difficult to resolve. A common method of contrast enhancement known as global histogram equalization is depicted in the right-hand image of Figure 2. Global histogram equalization is an automatic process that attempts to equally distribute the bins of the original histogram across the final rescaled bit range. When compared visually with the linearly rescaled image the equalized image contains more contrast and its histogram has been radically transformed. Flow induced temperature gradients at the leading edge of the wing, tail, and engine fairings are now evident. However, the fine spatial detail is lost in the oversaturation of the image.

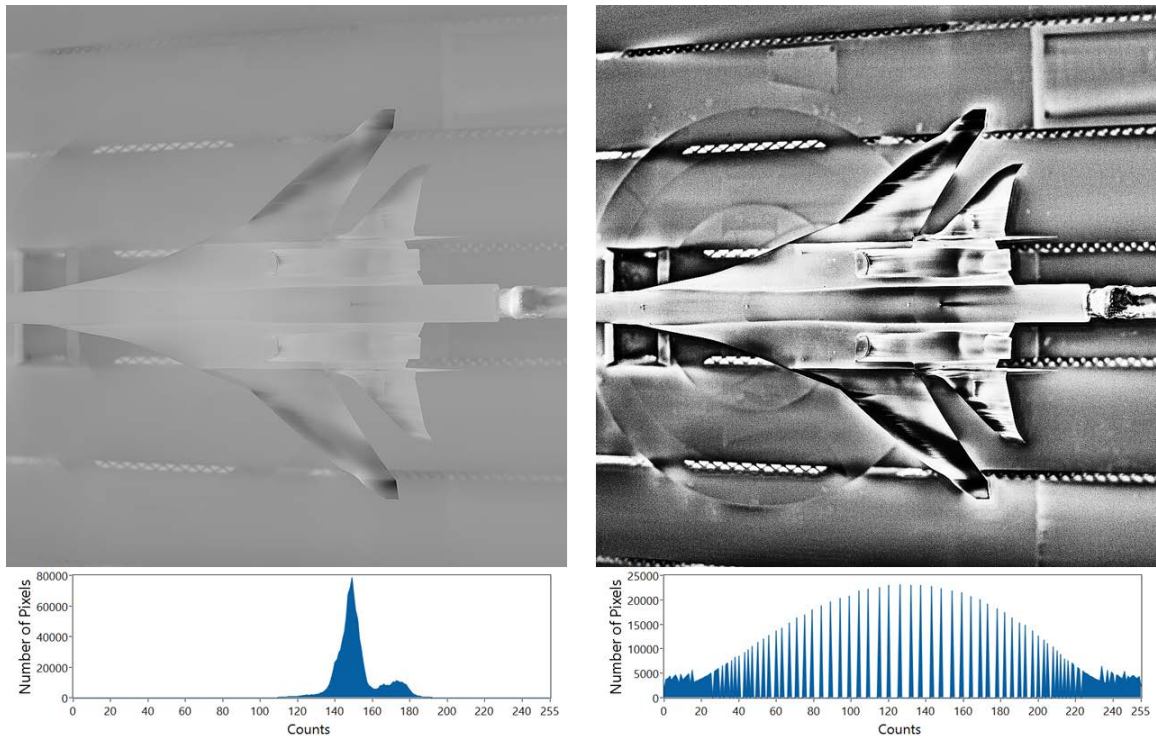


Figure 2: “Reference” IR image, 8-bit linear rescale (left) versus global histogram equalization (right)

A contrast enhancement technique known as adaptive histogram equalization (AHE) performs the histogram equalization operation within a fixed sized window that is convolved across the image. There are several implementations of this method, some reduce noise in uniform regions by limiting contrast, CLAHE, but most reduce compute time by equalizing within fixed sub regions and interpolating look-up-table values between [5,6,7]. With high dynamic range images, the interpolated AHE implementations can cause artifacts in regions with sharp pixel value gradients [1]. The implementation of AHE discussed here is non-interpolated (the equalization sub windows are convolved pixel by pixel across the entire image) to avoid these artifacts. Figure 3 compares two methods of AHE processing performed on the previously discussed reference image. The left-hand image has been processed with a sub-region of 128-by-128 pixels ($1/8 \times 1/8$ total image size). Qualitatively the contrast of the image has been improved while the small scale flow features have not been over-saturated. However, due to the fixed size of the sub-region the scale of the features that are enhanced do not exceed the size of the sub-region. Even with the non-interpolated AHE implementation image artifacts can still arise. Like the interpolated method this happens in areas of high gradients where the sub-region spans a large dynamic range. Often this is manifested as a light or dark halo effect around the edges of the test article. A slight halo effect can be seen in the left-hand image of Figure 3 near the leading edge of the wings. To address the shortcomings of a fixed scale AHE implementation a multi-scale AHE method is proposed, whose results can be seen in the right-hand image of Figure 3. In this multi-scale method two fixed regions are used to process the original image. One region is intended to enhance large-scale features using a large sub-region (in this case 256-by-256 pixels), while a second small region is used to enhance small-scale features (in this case 32-by-32 pixels). The result is two separately enhanced images, that are then blended to produce a single image. This multi-scale AHE method produces good quality results with minimal halo artifacts but is computationally expensive as the AHE method (computationally expensive in its own right) is performed twice.




Figure 3: Adaptive histogram equalization, single-scale (left) versus multi-scale (right)

Upon considering the single-scale and multi-scale AHE techniques it is apparent that choice of spatial scale is important. To that end a multiscale contrast enhancement approach will be considered. The Laplacian pyramid [8] is a method of decomposing an image into a multi-level set of spatially filtered band-pass images. Here it is proposed to extend this to a method of contrast enhancement based on the following: The histogram at each band-pass image level is generally normally distributed and centered about zero. Consequently, it is straight forward to normalize each image level based on the standard deviation of its histogram (Figure 4). Each level can then be weighted and the pyramid reconstructed to produce an enhanced image (Figure 5). This method seems to consistently produce a visually appealing image that is low in noise and able to simultaneously enhance many spatial levels of detail in the image (Figure 6). Additionally, the resulting image can be computed quickly as the decomposition is based on a set of separable filters that can be vectorized making this application ideal for a real-time system. The draw-back of this method is that it requires input from an operator to determine the normalization parameters and weightings of each spatial scale, whereas the multi-scale AHE method discussed previously is mostly parameter free.




Figure 4: Normalized Laplacian pyramid decomposition (six band-pass levels)





Figure 5: Normalized Laplacian pyramid reconstruction




Figure 6: Linear rescaling versus normalized Laplacian pyramid reconstruction ($\text{Mach} = 0.95$, $\alpha = 5^\circ$)

B. Dynamic Mode Decomposition and Discrete Time filtering

The following outlines a technique to use dynamic mode decomposition (DMD) and discrete time filtering to process infrared datasets. This technique will be shown to reveal low signal strength flow phenomena and is applicable to test articles with poor emissivity. DMD is a relatively recent tool originally developed to characterize complex fluid flows [9]. It is a data-driven method that combines features of principal component analysis with Fourier transforms in time. The results are physically interpretable spatial modes that are associated with a temporal frequency [10]. A key advantage of DMD with respect to traditional Fourier analysis is that DMD is able to capture the growth or decay of spatial structures in addition to their frequency of oscillation. This is valuable when analyzing non-stationary signals with transient behavior as is often the case in the context of IR visualization. In the proposed method, DMD is used as a diagnostic tool to identify flow structures in the data as well as their associated frequencies. These results are used as a guide to design a zero-phase bandpass filter which is then applied to the original time history data. The result is a dataset that has been filtered to enhance a particular flow structure of interest that is not apparent in the unfiltered dataset.

An example of this process is demonstrated in Figures 7 and 8. Infrared cameras were setup to image a high-aspect ratio aircraft in the NASA Ames UPWT 11-by-11 foot test section. Due to the nature of the wind tunnel model geometry the test article could not be painted with a high-emissivity coating. Typically this is a critical step in performing a successful infrared flow visualization measurement, as without such a coating the reflection of the surrounding test section dominate the image. A time-series of infrared images were recorded as the wind tunnel drive increased from zero speed to $\text{Mach} = 0.5$. DMD was performed on this data set and each mode examined for surface flow features. In this circumstance it was assumed the dominate flow feature would be boundary layer transition wedges so modes were sought out that reflected this phenomenon. Figure 7 details three selected modes extracted from the DMD solution. The images on the left and right of Figure 7 represent the modes corresponding to DC and 0.5 Hz respectively. These two images show no flow related temperature gradients on the surface of the wing. Like the unprocessed time series data, the dominating features are in fact reflections of the wind tunnel test section. These appear as diagonal lines and are projections of the tunnel slotted walls on the wing surface. The middle image of Figure 7 represents the mode of the same DMD solution set at 0.013Hz. This image distinguishes itself from the others in that there are flow features on the surface of the wing. The flow features take the form of transition wedges and are apparent out to ~ 0.06 Hz. This indicates that this narrow frequency range captures a boundary layer transition event as the Reynolds number of the wind tunnel increases. A zero-phase band pass filter is then applied to this frequency range to enhance the aforementioned transition event. The zero-phase property of the bandpass filter is important as it eliminates any phase distortion associated with the filter. The result is depicted in Figure 8, a boundary layer transition event is now apparent. The number of wedges increase and move forward towards the trip dot line located at ~ 33 percent chord as the Reynolds number increases. In summary the DMD results have indicated that boundary layer transition has occurred, while the tuned zero-phase bandpass filter details when the transition event occurs in the original time series. The measurement of boundary layer transition is a typical IR flow visualization result, what makes this novel is the fact that it has been imaged on an unpainted model in a wind tunnel test environment.



Figure 7: Three selected modes: DC, 0.013Hz, 0.5Hz (left, middle, right)




Figure 8: Boundary layer transition on unpainted steel wing-tip, filtered and unfiltered (top row, bottom row): $Re/ft = 0$, $1.43e6$, $2.45e6$ (left, middle, right)

Next, we examine a case where simultaneous IR imaging was performed on painted and bare metal wings. The previously discussed DMD analysis and filtering was performed on the bare metal image set while standard contrast enhancement techniques were used on the image set of the painted wing. This provided an opportunity to validate this method against a control dataset. Figures 9, 10, and 11 depict the onset of boundary layer transition as the speed of the wind tunnel is increased from no-flow to Mach 0.12. The left-hand image depicts the unpainted, left wing which has been processed using DMD and band-pass filtering. The right-hand image depicts the painted, right wing which has been processed using contrast enhancement techniques. The trip dot line was located near the leading edge of the wing and so a wedge transition front would be expected to coalesce at the leading edge. Figure 9 shows the wind-off images for reference. The surface of the unpainted wing is dominated by reflections of the test section. These largely appear as bright lines where the slots in the test section walls are projected onto the surface of the wing. The large, bright rectangles are back reflections of an IR camera. In contrast, the image of the painted wing demonstrates how effective a high-emissivity coating can be at removing these reflections. It should be noted that at this point in the wind tunnel test the painted wing had received a few chips in its finish. These appear as voids near the trailing edge of the wing-root junction, and an area at approximately two-thirds span. As the wind tunnel main drive is engaged, the flow begins to increase speed, boundary layer transition can be identified as turbulent wedges as depicted in Figure 10. The wedges on the painted and unpainted surfaces appear at nearly the same flow conditions, confirming that DMD based filtering has indeed positively resolved the boundary layer transition event in time. Finally, Figure 11 details the boundary layer as it nears coalescence at the trip dot line. These images are dominated by the finer structure of the transition front and can be seen in both unpainted and painted images. In conclusion, this technique seems to produce valid results for the case of time resolving boundary layer transition. It should be noted that these results can be difficult to interpret as the background reflection is typically not completely removed from the image data. This is likely due to the fact that some portion of the background reflections are changing at similar frequencies to the boundary layer transition event. Additionally, the authors are not advocating for the obsolescence of painted surfaces for IR flow visualization. In fact the painted surface improves the IR flow visualization due to both its high emissivity and thermal properties. While the

aforementioned DMD driven filtering technique can improve the results on low emissivity surfaces it cannot account for the loss in beneficial thermal properties of a painted surface; namely a lower thermal conductivity that helps establish the thermal gradient between the free-stream and the surface.




Figure 9: Simultaneous boundary layer transition, unpainted (left) / painted (right) wing, Mach = 0, $Re/ft = 0$




Figure 10: Simultaneous boundary layer transition, unpainted (left) / painted (right) wing, Mach = 0.07, $Re/ft = 0.6e^6$




Figure 11: Simultaneous boundary layer transition, unpainted (left) / painted (right) wing, Mach = 0.12, $Re/ft = 0.92e^6$

In the context of IR imaging, DMD may be used as an effective tool to reveal hidden structure in a dataset. This is largely due to the high bit-depth and sensitivity of modern IR camera sensors producing information rich data. Figure 12 demonstrates an unexpected result from IR image data compiled during a Mach sweep (1.1-1.4) in the wind tunnel. The results of the DMD revealed off-body terminal shocks on the horizontal tail of the wind tunnel model. Pass-band filtering of the dataset associated these terminal shock modes with Mach number. It is believed this is a sort of IR shadowgraph where the IR light source is the tunnel itself. This terminal shock flow feature is not visible in either the raw or contrast enhanced data. Its structure is only made apparent through the use of DMD and correlated in time by the application of a zero-phase passband filter.




Figure 12: Development of horizontal tail terminal shock: Mach = 1.1, 1.2, 1.4 (left, middle, right)

III. Conclusions

In the preceding discussion we have qualitatively explored the use of a few processing techniques to enhance IR flow visualization imagery. These methods were separated into two categories: contrast enhancement and a DMD driven discrete time filtering technique. Generally, contrast enhancement alone as a method of IR image enhancement is acceptable in circumstances when the test article being imaged possesses high-emissivity. A multiscale, non-interpolated adaptive histogram equalization technique was discussed and compared with an enhancement method based on a Laplacian pyramid reconstruction. The Laplacian pyramid reconstruction is less computationally intensive than the non-interpolated adaptive histogram equalization since it is based on separable filters. In the reference image studied here the pyramid reconstruction qualitatively produced superior results. It does however require more user input and may therefore be more difficult to implement in an automated fashion. In situations where the emissivity of the test article is low, we present a method of processing an IR image time series, which seeks to decouple the background reflections that dominate the unpainted surface from the aerodynamically driven surface temperature gradients based on temporal characteristics. This technique relies on DMD to produce a set of spatial modes each with a unique temporal frequency. In the case of IR flow visualization analysis, DMD possess an advantage over traditional Fourier based techniques in that DMD is well suited to analyze non-stationary signals. This DMD analysis step results in the identification of the structure and frequency content of the aerodynamically driven surface temperature gradients (if they are separable from the background). With this information in hand a zero-phase bandpass filter can be designed and applied to the original time series in order to establish the timing of the event in question. This method was used to positively identify boundary layer transition on an unpainted model in the 11-by-11 foot UPWT wind tunnel. Synchronized imaging on painted and unpainted wings during a boundary layer transition event verified the effectiveness of the zero-phase filter methodology to capture the event at the proper time within the dataset. Finally, this method was demonstrated to reveal hidden structures in the IR image data as shown by resolving the off-body terminal shock on the horizontal tail of a wind tunnel model in supersonic flow.

Acknowledgements

The authors would like to thank the following principal investigators for generous use of wind tunnel model imagery: Greg Gatlin (Truss-Braced Wing), Don Durston (Boeing N+2 Sonic Boom Performance model), and Max Amaya (Ames Check standard Model).

References

- [1] Garbeff, T.J., Baerney, J.K., “*Recent Advancements in the Infrared Flow Visualization System for the NASA Ames Unitary Plan Wind Tunnels*”, 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-1051)
- [2] Vollmer, M., Möllmann, K.-P., “*Infrared Thermal Imaging: Fundamentals, Research and Applications*”, Wiley-VCH, Germany, 2010
- [3] T. Acharya, A. Ray: *Image Processing: Principles and Applications*. Wiley-Interscience 2005, ISBN 0471719986
- [4] Garbeff, T.J., Heineck, J.T., et al, “*Institutional Schlieren: A Production-Level Wind Tunnel Test Measurement*” 53rd AIAA Aerospace Sciences Meeting, AIAA SciTech.
- [5] S. M. Pizer, E. P. Amburn, J. D. Austin, et al., “*Adaptive Histogram Equalization and Its Variations*” Computer Vision, Graphics, and Image Processing 39 (1987) 355-368.
- [6] D. J. Ketcham, R. W. Lowe & J. W. Weber, “*Image enhancement techniques for cockpit displays*” Tech. rep., Hughes Aircraft. 1974.
- [7] K. Leszczynski, S. Shalev: *A robust algorithm for contrast enhancement by local histogram modification*, Image and Vision Computing, Volume: 3 Issue: 3, 1989.
- [8] Burt, P.J., Adelson, E.H, “*The Laplacian Pyramid as a Compact Image Code.*” IEEE Trans. Commun.(1983) 532–540.
- [9] Schmid, P.J., *Dynamic Mode Decomposition of Numerical and Experimental Data*. Journal of Fluid Mechanics, 5-656:5-28, August 2010.
- [10] Kutz, J.N., et al, *Dynamic Mode Decomposition Data-Driven Modeling of Complex Systems*. Society for Industrial and Applied Mathematics 2016, ISBN 9781611974492